Financial Model

for Measuring Cost Savings Driven by Service Operations that Impact Customer Experience and Voluntary Disconnects

By Robert F. Cruickshank III, Chief Executive Officer, Chief Technology Officer, Rev2

This article describes an automated, systemic solution for identifying hidden issues using existing data from the service delivery infrastructure - before they snowball into costly customer-impacting issues and voluntary disconnects.

Overview

Cable MSOs generally detect outages in node-serving areas by looking for real-time threshold violations in the number of recent subscriber "connectivity" phone calls, truck rolls and MTA/modem/set-top box de-registrations. For example, a call centre will typically declare an outage with call volume at or above a threshold of approximately three calls per node per hour. Yet there are few automated solutions for automatically processing a low-level of "dribbling in" of calls and truck rolls over several days or weeks.

Likewise, while outage detection exists for multiple customer premises equipment (CPE) devices falling offline (i.e. number of offline devices rising above a certain threshold), there are few automated solutions for smaller deviations in offline devices, and insufficient linkages among disparate database records of subscriber calls, truck rolls and offline devices.

In short, with real-time analysis, small and intermittent issues/outages may not exceed thresholds. Compounding the problem is that issues may be hidden for extended periods. Further, some issues are precursors of larger outages.

This article describes an automated, systemic solution for identifying these hidden issues using existing data from the service delivery infrastructure – before they snowball into costly customer-impacting issues and voluntary disconnects.

In addition, the article describes a methodology for prioritising maintenance of these issues according to their cost and their potential impact on the customer experience, as well as a financial model for calculating operating cost savings generated by the methodology.

> **"There are few automated solutions for automatically processing a low-level of "dribbling in" of calls and truck rolls over several days or weeks."**

Identifying the sweet spot

To begin, let's take a look at the typical costs of service and maintenance activities at a cable operator. Figure 1 depicts a preventative maintenance model that shows the ideal financial position for a cable operator (at the centre of the graph).

The target at the centre represents balance. It is the point of equalisation between the cost of customer service and field service on the left and the cost of maintenance on the right. If we perform very little maintenance on the network as depicted on the left-hand side, we would logically have very high calls and truck rolls as a result. On the right-hand side, if we were to perform a very high amount of maintenance, there would probably be diminished call volume.

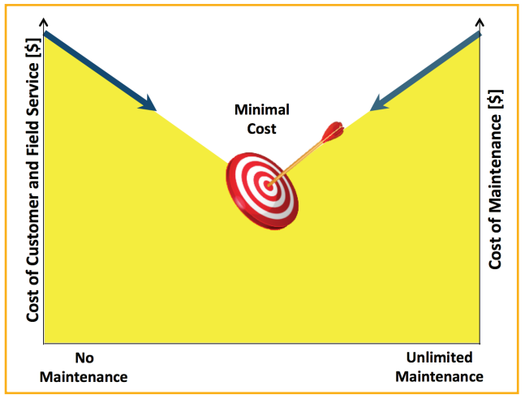

Figure 1. Levels of preventative maintenance. The goal of the cable operator is to minimize the cost of both service and maintenance at the same time.

So with our methodology, our goal is to help the operator find the balance or the sweet spot in terms of spending between maintenance and service activity plus call volume activity whilst simultaneously minimizing voluntary disconnects.

Problem resolution flow

To help understand how we can help operators achieve this balance, let's look at Figure 2. This flow chart depicts the sequence of events that occurs when a customer calls with a valid HFC connectivity issue that has yet to be detected by the business. There is a problem in the plant, which is affecting the end user experience to the point that customers are calling. When there is a problem in the service delivery infrastructure, it manifests with customers calling about a number of connectivity issues.

Some typical examples include but are not limited to:

- Voice issues, such as no dial tone, intermittent dial tone or low quality of service.

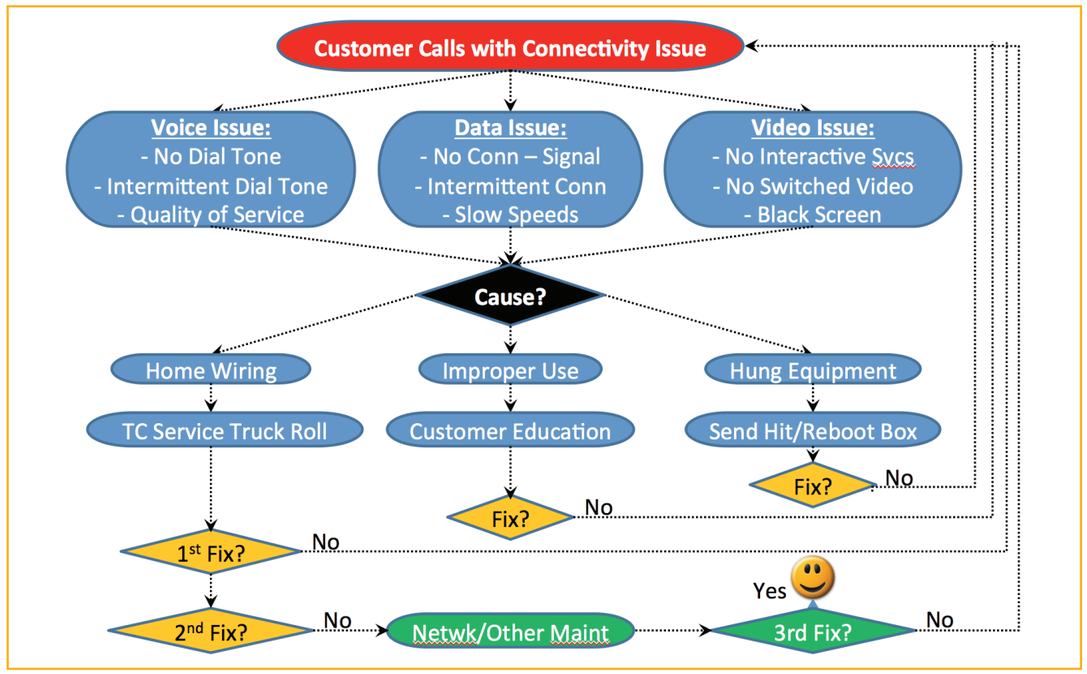

Figure 2. Call centre, service and maintenance life cycle of subscribers with valid HFC connectivity issues

- Data issues, for example they can't connect at all, they can only connect intermittently, or they can't download or stream.
- Video issues, such as they can't order two-way services, they can't be interactive with the network, they can't order channels, they have a black screen or they have pixellation, freeze frames or audio drop-outs.

Each of these issues can be manifestations of connectivity problems that may actually reside in the plant somewhere. So if there is bad ingress, bad power levels or high error rates in the HFC infrastructure, the subscriber could experience any of these issues.

Let's look at the flow of the customer call in Figure 2 in sequence. Across the second row, the call centre service representative is faced with the challenge of categorizing each call based solely on the customer's report. It is straightforward for the service rep to categorize a call in the area of voice, data or video. But there is no way for them to recognize that these issues might be related. We will look at how calls can be correlated with field service/maintenance activity and service delivery infrastructure information to establish a common correlation between events.

Guessing the cause

In the diamond-shaped box called "Cause", the customer service rep has to – as quickly as possible and based only on information from a telephone call – determine what he/she thinks is the cause of the problem. On the right-hand side of Figure 2, customer service has decided that it needs to re-set the equipment because it is hung (not responding) or needs to have services re-authorized/refreshed. To remedy this, it can send a hit to the box or reboot the modem, or perform any number of combinations of diagnostics.

Sometimes this remedy will make the box come back online and provide service again. Both the subscriber and the customer service representative are hoping this repair is going to solve the problem long-term. However, if this problem is actually in the plant and not in the modem, it will rear its ugly head again – not really a long-term fix. As a result, you get into situations where the customer calls back again.

Another possible resolution is in the middle path. The service rep could determine that the modem has been improperly used and, after another reboot, could recommend a resolution around customer education. That could include teaching the customer to reboot the modem on his own, how to clean out the cache on their browser, or even how to troubleshoot some of their wiring. At the end of the call, the resolution will be coded as customer education. This is the highest resolution category that we've seen.

But if we go back to this assumption that the problem is really in the plant, no kind of customer education is going to fix the problem either. The customer will eventually call back again and again so there will be no resolution. At some point, we would go down the left-hand path where the service rep decides that the problem is a home wiring issue. This is the path taken after remote troubleshooting has failed and it is necessary to send a truck to the home. When the maintenance person visits the home, they will probably perform one or more of the following five typical resolutions:

- Replace the modem or set-top box
- Run a new drop
- Run new inside wiring
- Change a splitter configuration, or
- Troubleshoot the ground block.

Financial cost and customer experience

Coming back to the assumption that the problem was really in the plant, then of course none of these actions solved the problem. Each of the three paths is increasingly expensive to the operator, with the truck roll being the most expensive. In addition, when unsuccessful over time they are increasingly damaging to the customer experience.

The expense is compounded for the operator when a high percentage of equipment – set-top boxes, modems etc. – are returned, often with "No Trouble Found" because the problem is in the plant but nobody has identified it yet. This begs the question, why hasn't the problem been identified yet?

The answer is that the problem wasn't big enough to be observed. It may not have crossed a pre-existing real-time threshold. For example, it didn't cause a large enough hard outage. As a result, the customer will continue having difficulty and will probably call back again; he or she will most certainly call back if they experience a hard outage. The result is ultimately a second or repeat truck roll.

The customer has already been a repeat caller and now he requires a repeat truck roll. It is at this point that the service

"A typical operator generates so many alarms that engineers tend to ignore all but the most critical issues."

department representative may decide the problem isn't in the home but in the plant, and then he may escalate the issue to the plant maintenance department. When that happens, it's really the first time that someone who has been empowered to fix the actual cause of the problem has been involved in the situation.

Maintenance teams in the plant are constantly scanning the service delivery infrastructure for problems. So if an issue is hidden from their detection, they will find out about it from this escalation. This is generally known as a "refer to maintenance" – where the service department has asked plant maintenance to step in – and again it is costly to the operator and detrimental to the customer's experience.

As you can imagine, by the time a "refer to maintenance" occurs, the customer is distraught, frustrated and angry, and thinking that it's just too difficult to do business with this cable company. At this point, the customer experience is off the charts and the customer may in fact choose to disconnect.

Why problems are hidden

A challenge that operators often face is the "silo" nature of operational databases that store outage risk data. There is often one database that contains troubleshooting records from voice and data subscribers; a second database with troubleshooting records from video subscribers; a third database with truck rolls to subscribers; a fourth database with physical plant maintenance truck rolls; a fifth database with network telemetry readings etc. Generally, these databases are dissimilar enough that aggregate analysis of their data is time-consuming and tedious.

Because of these challenges, smaller outages and pockets of degraded service may go undetected long enough for repeat calls to manifest as complaints to executive management – often resulting in staff ultimately finding and validating an actual subscriber-affecting issue and then regretfully agreeing "Why didn't we see that earlier?"

A typical operator generates so many alarms that engineers tend to ignore all but the most critical issues. By correlating, classifying and aggregating micro alarms, the operator is provided with a very high probability of detection and a very low missed-issue rate, alleviating a major drawback of current alarm technology.

Automating problem finding

To facilitate the early identification of problems in the plant, an automated system keeps track of all transactions – calls to the call centre, maintenance truck rolls and escalations to the plant. The system captures these problems systemically and correlates them to raise flags on issues faster and earlier than they would have otherwise been detected. This methodology is enabling operators to get ahead of calls to the call centre and save costs on service and maintenance.

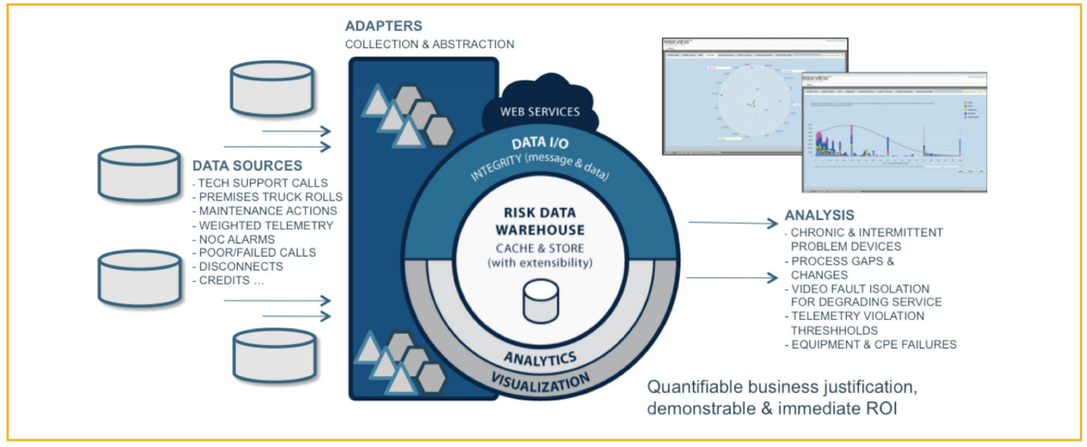

Figure 3. Data warehouse input/output model

Real-time analysis thresholds are triggered after three calls in an hour to the call centre, or by a pre-established number of modems going offline. It's clear that these thresholds can't be set low enough to find the smaller problems which are often the precursors to bigger problems. Of course these thresholds are set so that maintenance and service organisations aren't unnecessarily loaded with work – especially when that work will undoubtedly result in many "No Trouble Founds".

Unfortunately, intermittent issues and outages will be hidden for an extended period of time – until they get worse. And they usually do get worse. Fortunately, the good news is that, thanks to our service and maintenance calls, we have early warning clues that can lead operations teams to concentrations of risk underlying the service delivery infrastructure in the plant that have not yet been recognized by the "business as usual" procedures. Those clues are needles in haystacks because there's a lot of data around them and they are imprisoned in these disparate databases. With the data warehousing methodology shown in Figure 3, we've found a way to unlock this and identify problems early based on Risk Concentration Analysis.

Risk concentration analysis

Multiple records are used to point to material concentrations of risk. A critical advantage of the risk concentration analysis methodology is classifying risk data in a meaningful way so that operators can see these concentrations. These "risks that matter" become evident when each risk is considered in the context of all other risks existing throughout the service delivery and support infrastructure. The methodology addresses the reality that different risks have different impacts on the business. The primary differentiator is determining where there are concentrations of risk. With risk concentration analysis, material risks emerge when correlating risks from all silos and considering each risk in context of impact to the business. This approach yields Financial and Reputational scores, making it easy to recognize and prioritize material risks.

This is accomplished by collecting all risk data from across the operator's business and service delivery infrastructure, and then normalizing risks into a common format and language so that they can be compared by assigning a unique Financial and Reputational score to each risk.

The score of each risk reflects its materiality to the business or subscriber – as well as the impact that problem would cause. When visualized, the risks that represent the greatest vulnerabilities will stand out from the rest. This helps operators to identify risks and then test controls in the context of all other risks, as opposed to looking at risks in isolation. This identifies, for example:

- Chronically misbehaving devices
- Recurring problems specific to a geographic region or departmental silo
- Problems in programme execution that are impacting the business's reputation or financial bottom line – and could escalate.

Data sources and analytics

As you can see on the left-hand side of Figure 3, the sources for the data warehouse engine are calls to the call centre including trouble calls and the disconnects associated with them; service reports; dispatch to field operations and network operations centre data on failed network telemetry. In this case, the ARRIS WorkAssure® provided the TC service calls and repeat TC data, along with the EC ("refer to maintenance") activity. These metrics are then correlated to look for commonalities, such as shared regions, hubs, nodes or problem IDs.

In addition, the data warehouse algorithm adds weight depending on whether a customer is a business or residential customer or whether they subscribe to one, two or three services. In this methodology, we only use maintenance activity that is in response to customer-reported service delivery infrastructure issues. For example, if maintenance has been working on the plant, we can look back to see if that work was successful or had to be repeated. We can also choose to not include maintenance activity that was a planned event, such as capacity addition. However, we can look back to make sure that work was performed properly so that all the performance levels came back to the same level they were before.

> **"The methodology addresses the reality that different risks have different impacts on the business."**

The final data input is metrics from telemetry systems, whether they are home grown, off-the-shelf systems or some combination of both. In this example, we're using the ARRIS ServAssure® system which is telling us about the number of modems that are in trouble every hour by reporting upstream and downstream errors, noise, power and online and offline status. And we're also able to look at combinations of nodes, also known as "serving groups".

As shown in the centre of the diagram, the data warehouse architecture collects the data sources via adapters. Adapters are simple PERL scripts that normalize all of the collected data into a common format and apply the assigned weights for correlation. Typically, results are reported to the Operations team once a day – say 6:00 am at the beginning of the day shift. We've found that it's getting easier to redeploy this same architecture and we've done it with several operators to date.

The role of DOCSIS telemetry

As mentioned, another insightful source of information is telemetry data from in-home and in-business CPE devices such as cable modems and set-top boxes that support the DOCSIS (Data Over Cable Service Interface Specification). In concert with network elements such as cable modem termination systems, DOCSIS devices provide remote access to several metrics such as Uncorrectable Codeword Error Rates (CER), elevated re-set behaviour and unusual online/offline behaviour – on both the shared downstream and upstream channels as well as to and from each individual device.

Even more insightful, depending on the frequency and magnitude of DOCSIS telemetry readings, the level of subscriber pain can vary. Examples include:

- A cable modem that spontaneously re-sets once a week is less of a problem for a subscriber than a modem that resets tens or hundreds of times per day.
- A cable modem with -9 dBmV downstream receive power is less problematic on hot summer days than a modem with -9 dBmV CMTS upstream receive power.
- A cable modem or set of cable modems with high Uncorrectable Errors is more problematic than a similar set of modems with high Correctable Errors. And these are of greater concern than a set of modems with very low numbers of Correctable Errors.
- An important requirement is uniquely classifying and scoring DOCSIS CPE telemetry data so that those specific telemetry readings having greater impact to subscribers are given greater importance and higher Reputational Cost scoring. By doing so, an aggregate tonnage of risk concentration can be easily calculated and then used to prioritize maintenance and repair efforts.

Figure 4. Outage risk analytics radar chart showing costs of nodes by region. The outliers represent the most expensive nodes and can be used to prioritize maintenance issues.

Financial cost case study

Now let's look at the quantifiable business justification for the system. The service outage analytics are presented via data visualization. The cost of operation is shown via a radar chart as shown in Figure 4. Each dot represents whatever makes the most sense to the operations analyst i.e. it could be a node, a serving group, a hub or a street. The distance each dot lies from the centre of the diagram is based on cost. In this case, the outliers represent the most incurred cost (already out the door) to operate the various segments of the network. So what does this tell us when we get down to a detailed level in the real world? Let's look an example study that we conducted at a major North American MSO in 2012. We conducted a blind comparison between business as usual at the MSO (the Control Group) and compared it with issues that were found with the Risk Concentration Analysis methodology (the Experimental Group).

The baseline statistics of the study area were as follows:

- 1,000,000 subscribers
- 25,000 miles of plant
- 180 technicians, 15 supervisors
- 7,500 nodes in 3,000 serving groups.

The scientific hypothesis that we set out to prove was that operating without the Risk Concentration Analysis methodology was going to be greater than operating costs with the RCA automation. We also wanted to prove via straightforward mathematics that the resulting difference would be a quantifiable saving. As you'll see, the hypothesis was proven.

The Control Group inputs from 'business as usual' included connectivity trouble calls, truck rolls, telemetry readings and maintenance activities. The Experimental Group used the same inputs and deployed the RCA correlation engine in an automated fashion to conduct outage risk analytics. The outputs were directed towards two groups:

- The service department, which looked at non-area issues, and
- The maintenance department, which looked at area issues.

Figure 5 shows our Cost Savings calculation method. After correlating the costs across the different DOCSIS serving groups, this equation was used to alert the two groups on unusual costs. In this calculation, we weighted each input to allow us to look at Reputational Costs as well as Financial Costs. The Reputational Costs include things that can directly impact the customer's experience: people calling in, trucks going out, repeat trucks, escalations to maintenance, disconnects, demand and preventative maintenance activities, upstream and downstream errors, noise, power, re-sets etc.

In addition, we created an assumption that by reducing customer issues during that time period, we can reduce 25% of the final operating costs. In this case, we used the cable operator's actual costs: 55 cents per minute for Calls, US$ 60 for truck rolls (TC) and Disconnect (DI), and US$ 100 for

Base Weights Applied

1. Reputational Cost was used to trigger plant maintenance alerts:
Reputational Cost = Calls + (2*TCs) + RepTC + (3*ECs) + (2*DI) + CHG + TRB + Up&Down (Err + Signal Noise Ratio + Power + Resets)

2. During alert periods, 25% of Financial Operating Cost was used to calculate savings:
Financial Operating Cost = US$ 0.55/min Calls, US$ 60 TC and DI, US$ 100 EC, CHG, TRB
Note: costs provided by MSO finance director.

Alert Period

The Alert Period = Prior Day Reputational Cost ≥ 15
AND (Same Day OR Next Day Reputational Cost ≥10)

Total Operating Cost US$ 3,228 =
71 Calls + 24 TCs +2 EC + 12 DI + 1 CHG + 2 TRB

Potential Savings based on Alerts driving earlier maintenance actions:
US$ 411 = 7.5 Calls + 2.75 TCs + 0.25 EC + 1.5 DI + 0.25 CHG + 0.5 TRB

Figure 5. Cost savings model calculations.

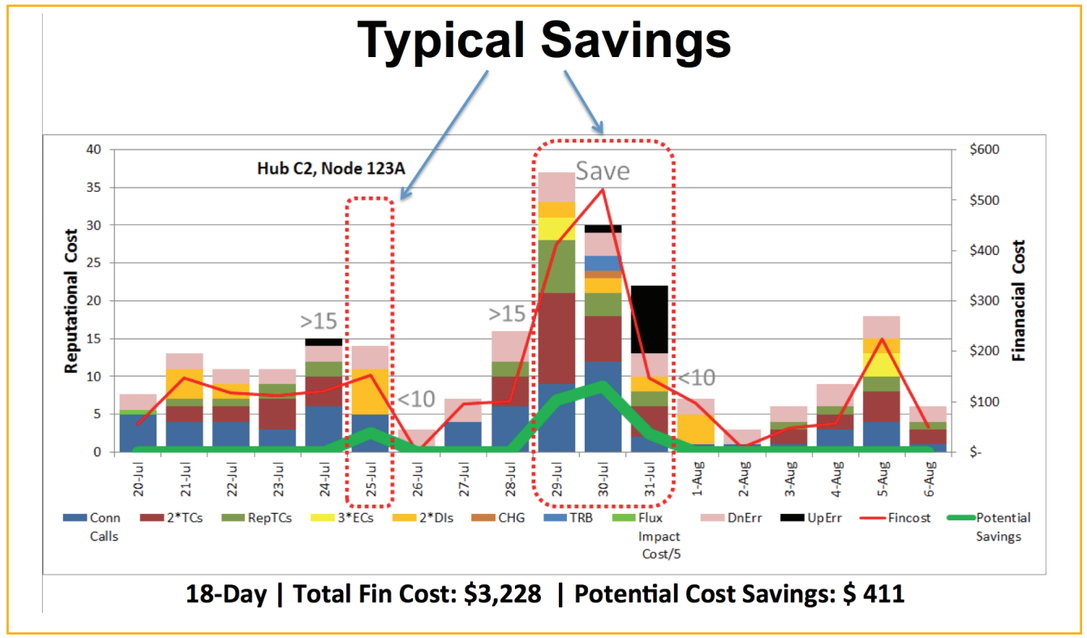

Figure 6. This Outage Risk Analytics financial model shows connectivity issues underlying a typical serving group, node and street behaviour over an 18-day period. Two periods of potential savings are identified by red dotted lines.

escalations to maintenance (EC), change (CHG) and trouble tickets (TRB). If we were to look at a single node and we wanted to calculate the savings on that node, we would have an Alert Period, also detailed in Figure 5.

The Alert Period is a period of heavy issue response and activity during which we can start calculating savings. The Alert Period is calculated by an on-trigger and an off-trigger. In this study, we determined the periods of the most costly activities spiked when the reputational cost was greater than or equal to 15. Therefore, 15 was determined to be the on-trigger.

As long as the days that followed sustained a reputational cost greater than or equal to 10, we determined that during that period we can capture Savings, and we can also capture Operating Cost, with Savings being a portion of the Operating Cost in that region. Then we multiply costs by the number of service calls and truck rolls etc. to get to the Cost Savings.

In the case study, we broke each serving group (groups of nodes) down into transactions. In Figure 6, the transactions are represented by the colours in the stacked bars. The left axis is the reputational cost axis. Starting at the lower left in blue, we received five connectivity calls that came in on 20 July. Starting on the second day, 21 July, we have the emergence of a couple of trouble calls (TCs), shown in red.

We have downstream errors in pink and upstream errors in black. We have escalations to maintenance (ECs) in yellow, some Repeat Trouble Calls in green etc. Our mathematically-calculated Alert Period, also known as the Savings Window,

Customer Type	ARPU	Churn	CLV
Analog	$45.00	2.50%, 40 Months	$222
Digital	$67.00	3.00%, 33 Months	$599
Video/Data	$100.50	2.00%, 50 Months	$2057
Video/Phone	$93.50	2.00%, 50 Months	$2062
Data/Phone	$71.00	2.00%, 50 Months	$1790
Triple Play	$133	1.00%, 100 Months	$5642
Triple Play	$133	2.00%, 50 Months	$2972

Figure 7. A financial model showing the Customer Lifetime Value of a cable subscriber. CLV = Present Value of (ARPU – COGS - Care Costs). Lower Care = Higher CLV.
Source: Dr. Ron Rizzuto, The NCTA Cable Show, May 2012, Boston, Mass.

is shown with the red dotted line. In this example, the first savings window lasts a period of one day and the second window lasts for three days. The rising financial cost, reflected on the right axis, is traced by the red line for each day.

As explained earlier, during the savings period, we've made the assumption that, by identifying problems early through Risk Concentration Analysts, 25% of the financial cost would be Savings. This assumption means that 25% of the transactions in that time frame were not counted in the Experimental Group's model because they would have been avoided. We assert that by getting to those issues before they impact more subscribers, we can cut out much of the issue response activity which otherwise would end up costing the company money.

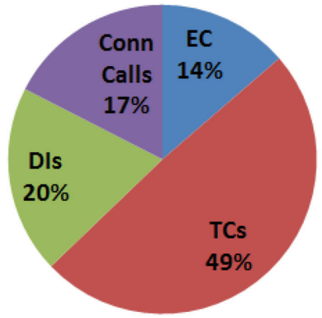

Figure 8. Pie chart showing percentages of each type of connectivity issue.

Customer Lifetime Value (CLV)

Next steps to include in the model include a calculation for lingering Issues – things that didn't cross threshold (yet) but still are causing significant pain. We are able to lower the Trigger Alert even further, based on accumulated cost over several days before it triggers. Based on all of these metrics, we can then assign a dollar value to the customer experience, according to the Customer Lifetime Value (CLV) formula developed by Dr. Ron Rizzuto at the University of Denver.

According to Dr. Rizzuto's formula, we know that different customers churn at different rates and that there is a different average revenue per unit (ARPU) on a monthly basis. Figure 7 shows Dr. Rizzuto's model. You can see that a current analogue customer has a lifetime value in the US$ 200 range, whereas a triple-play customer can have a CLV of US$ 3,000 to more than US$ 5,500 depending on their churn rate. And in the two triple-play calculations, you see the mathematical influence of churn and the difference it makes to CLV.

The important lesson here is that by trending down on transactions, we can also cut down on customer disconnects. Most importantly, when decreasing transactions, customer satisfaction will go up and churn will go down. In short, using the Risk Concentration Methodology, you can deliver a lower customer care cost and a higher customer lifetime value in total.

In the final breakdown of transactions, shown in Figure 8, it was learned that Service Trouble Calls (i.e. truck rolls) represented the most instances (49%), and Connectivity Calls to the Service Call Centre was second (17%).

Maintenance Department: Additional Work

16,000	Preventative maintenance actions created
-4,000	Saved escalations to maintenance/year
= 12,000	Net new preventative maintenance actions/year
US$ 1,200,000	Cost due to additional tasks @ US$ 100 each

Service Department: Reduce Contractor Head Count

Approximately 50% of installs performed by contractors @ US$ 60.00
As TCs are reduced, staff do "would be contracted" installs

US$ 900,000	Saved 15,000 TCs/year @ US$ 60
US$ 360,000	Saved 6,000 Disconnects/year @ US$ 60

Call Centre: Reduce capacity with normal staff attrition

US$ 320,000	Saved 40,000 Calls/year @ US$ 8

Projected Annual Savings:

US$ 380,000 net per year for 1,000,000 subscribers
= US$ 0.35 to US$ 0.40 per subscriber per year.

Figure 9. Additional cost savings metrics.

The final savings calculation

Before we obtain a final savings calculation, for balance we also have to consider additional work that was created by introducing new maintenance alerts into the business as usual procedures. Over the course of the blind study, we had 16,000 Savings periods during which we recommended a re-prioritizing of work by discovering problems early, as shown in Figure 9.

Using our 25% assumption for our Savings calculation, the Saved tasks represent about 4,000 escalations to maintenance. This leaves us with 12,000 net new preventative maintenance actions per year. At a cost of about US$ 100 each, that's a savings of US$ 1.2 million.

On the Savings side, one of the key areas of financial savings is reducing the number of overall tasks and therefore the manpower required by expensive service contractors. For example, in the study there were Savings on contractors because installs could be performed by in-house staff.

In the study, we calculated that using the RCA methodology, the operator can save up to 15,000 TCs, up to 6,000 Customer Disconnects (using the 25% assumption) as well as reduce calls to the call centre. When you add that up and subtract the US$ 1.2 million additional costs, the result is a US$ 380,000 savings per one million subscribers, or about 38 cents per subscriber per year.

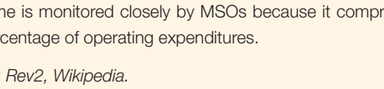

Abbreviations and Acronyms

CC – Call Centre. An office operated by a cable company to administer telephone-based support and information inquiries from subscribers.

CPE – Customer Premises Equipment refers to equipment located at a subscriber's premises that are connected with a carrier's telecommunications channels. This generally includes devices such as telephones, routers, switches, set-top boxes, fixed mobile convergence products, home networking adapters and Internet gateways that enable subscribers' access to services from the home.

DOCSIS – Data Over Cable Service Internet Specification. An international telecommunications standard that enables the addition of high-speed data transfer to an existing CATV system. It is employed by many Cable MSOs to provide Internet access over their existing infrastructure.

Financial Risk - In the Cable MSO world, financial risk is calculated using metrics such as truck rolls, call centre calls and churn rate.

Network Telemetry - A technology that allows remote measurement and reporting of information. Although the term commonly refers to wireless data transfer mechanisms (e.g. radio), it also encompasses data transferred over other media, such as a telephone or computer network, optical link or other wired communications.

NOC – Network Operations centre, pronounced "knock," a NOC is one or more locations from which control is exercised over a computer, television broadcast or telecommunications network.

NTF – No Trouble Found. A term used in various fields, especially in electronics, referring to a system or component that has been identified for repair but operates properly when tested. This situation is also referred to as No Defect Found (NDF) and No Fault Found (NFF).

NPV – Net Present Value. In finance, the Net Present Value of a time series of cash flows, both incoming and outgoing, is defined as the sum of the Present Values (PVs) of the individual cash flows of the same entity.

Outage Risk - The likelihood that a service will be disrupted at some point during its transmission, preventing it from being delivered to its destination subscriber.

RCA – Risk Concentration Analysis. Material risks emerge when correlating risks from all silos and considering each in context. The RCA approach yields a Materiality Score, making it easy to recognize and prioritize material risks – the risks that matter. After collecting vulnerability data throughout the business, RCA software assigns a score to each "risklet" that reflects its likelihood to cause a problem – as well as the impact that problem would cause.

Reputational Risk - Reputational risk is related to the trustworthiness of the business. Damage to a firm's reputation can result in lost revenue or destruction of shareholder value, even if the company is not at fault. Metrics used to calculate reputational risk in the Cable MSO world include the number of subscribers impacted, the number of services impacted and the number of outage minutes.

Truck Rolls - Refers to the act of dispatching a cable truck to resolve a service problem, usually at a home or street location. Truck roll volume is monitored closely by MSOs because it comprises a large percentage of operating expenditures.

Sources: Rev2, Wikipedia.

Bibliography

1. *"Drill Here: Emerging MSO Technology: Systemic Prioritization of Undetected Service Outages to Improve Overall Business Performance,"* a white paper prepared for the Society of Cable Telecommunications Engineers by Robert F. Cruickshank III, Chief Technology Officer, Rev2 and Matt Bell Managing Director, Teambell Consulting.